



epcc



National
e-Science
Centre

GridWeaver

Investigation of Grid & Utility Computing Fabric Management

NeSC Review

Edinburgh, 29 November 2002

Who are we?

- GridWeaver is a project run by:
 - School of Informatics, University of Edinburgh (*SoI*)
 - HP Laboratories, Bristol (*HP Labs*)
 - EPCC, University of Edinburgh (*EPCC*)
- The common denominator is active interest and research in the configuration of large installations of computing resources
 - *Fabrics*

What is Configuration?

- An effective Grid assumes the existence of correctly operating large-scale fabrics, on which the middleware and applications can be hosted.
 - What happens if one node in 500 has the wrong version of a math library?
 - What happens if we are unable to deploy a security patch quickly enough?
- We define *configuration* to include:
 - Initial software installation
 - Initial customisation of the nodes
 - Node reconfiguration to implement changes in the specified configuration
 - Monitoring and correction of unwanted changes in the actual node configuration

Configuration tools

- Traditional cluster configuration tools tend to be based on *imaging* (aka *cloning*) procedures.
 - A manual *golden image* is duplicated with slight changes to many other nodes
- Although these tools can scale effectively to very large clusters, they break down in the face of:
 - Diversity
 - Rapidly changing configurations
- Most large sites have found the need for tools which address these issues, such as those developed specifically to support workstation configuration:
 - CERN (SUE)
 - Argonne (Sanity)

Problems

- Tools such as *LCFG* from Edinburgh probably represent the state-of-the-art in workstation configuration technology.
 - The EDG testbeds currently use LCFG as a prototype
 - EDG WP4 are now designing their own tools based on the LCFG architecture, and these will have a similar functionality
- We have identified a number of problems with current tools which will become more serious as fabrics increase in complexity:
 - Specifying complex configurations and relationships clearly
 - Supporting devolved management
 - Sequencing configuration changes
 - ... and other issues as described in the project report

LCFG and SmartFrog

- We believe the way forward for very large scale configuration requires:
 - Models of the systems being configured
 - Languages to describe the desired configuration of the models
 - Systems to deploy the specified configurations on the target hardware
 - Systems to monitor and correct discrepancies between the specified configuration and the actual configuration
- *SmartFrog* from HP Labs includes a specification language and a deployment engine for distributed applications
- The aim of the project is to use the combined experience of LCFG and SmartFrog to investigate solutions to some of the above problems.

Organisation

○ PMB

- Paul Anderson, Sol: Project leader; Joint technical leader
- Peter Toft, HP Labs: Joint technical leader
- Kostas Kavoussanakis, EPCC: Tracking

○ Technical Staff

- Guillaume Mecheneau, HP Labs
- George Beckett, EPCC
- Carwyn Edwards, Sol
- Jim Paterson, Sol (from January onwards)

○ Technical Advisory

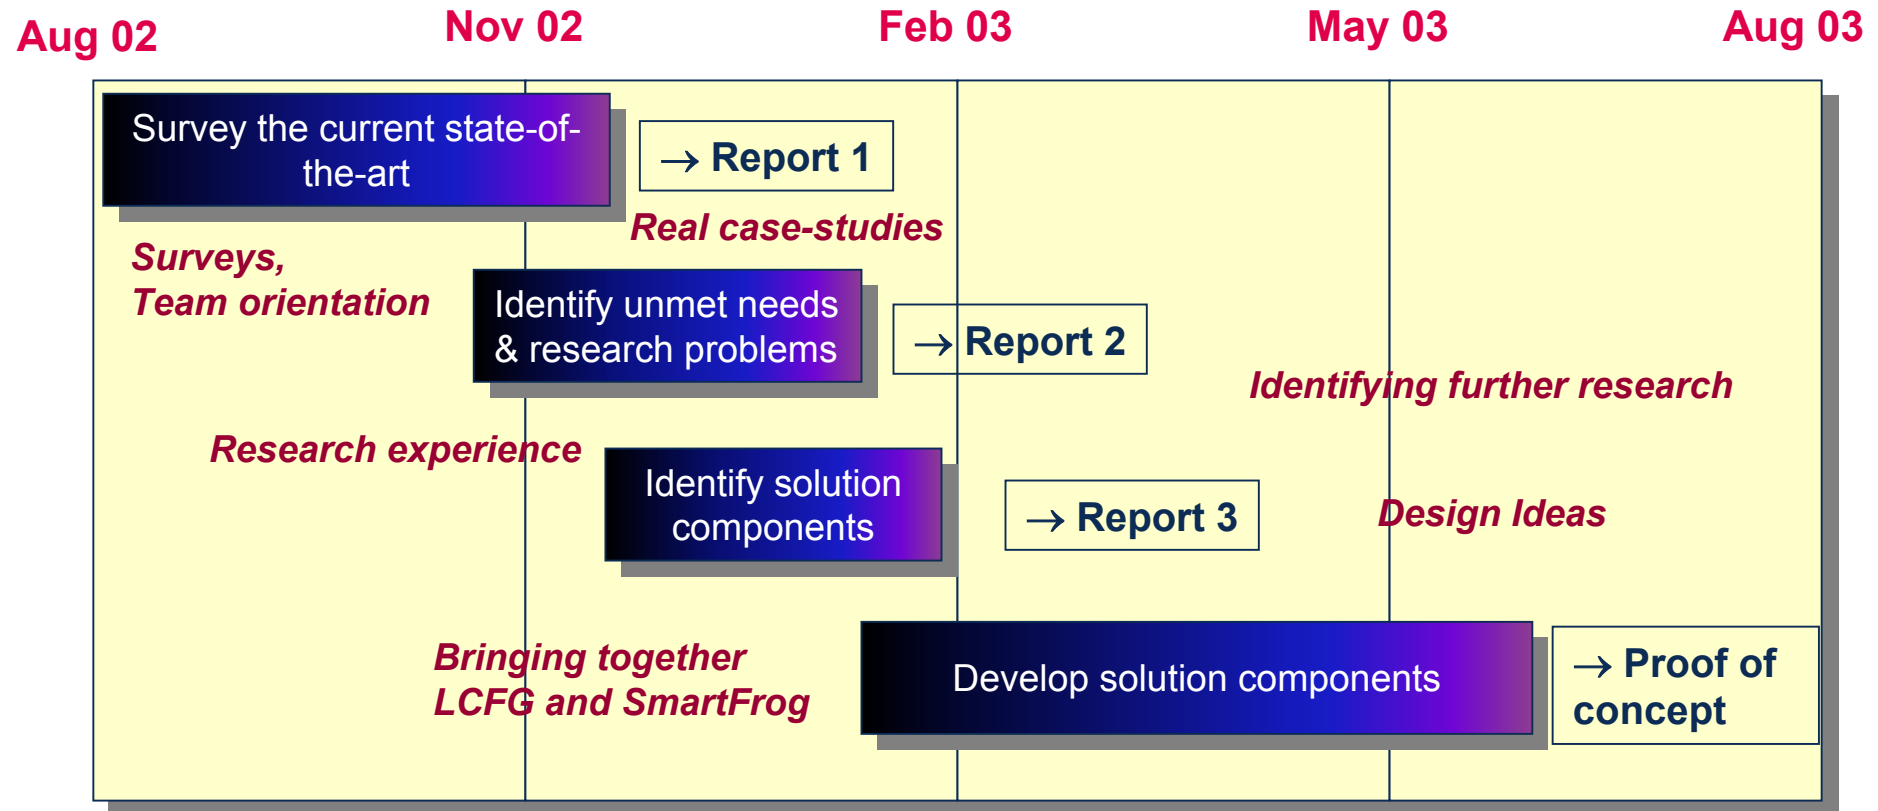
- Patrick Goldsack, HP Labs
- Alexander Holt, Sol

Modus Operandi

- Face-to-face workshops approx. every 6 weeks, alternating between Edinburgh and Bristol
 - Technical discussions
 - PMB issues
- Collaborative working supported by
 - A secure, shared repository
 - HP's *eRoom*
 - Extensive use of email, mailing lists and archives
 - Common CVS repository will be used for shared development

Workplan

- Phased workplan, with overlapping phases



D1.1: Technology Survey

- First deliverable, nearing completion →
- Sets out the principles of Large-Scale Configuration Management
- Surveys a range of representative commercial and non-commercial technologies
- Draws overall conclusions about the key approaches in current use
- Includes glossary of common configuration terms

Technologies for Large-Scale Configuration Management

Paul Anderson <paul@dcs.ed.ac.uk>
George Beckett <george@epcc.ed.ac.uk>
Carwyn Edwards <cedward1@dcs.ed.ac.uk>
Patrick Goldsack <patrick.goldsack@hplb.hpl.hp.com>
Alexander Holt <alex@fixedpoint.org>
Kostas Kavoussanakis <kavousan@epcc.ed.ac.uk>
Guillaume Mechenau <guimec@hplb.hpl.hp.com>
Peter Toft <peter.toft@hp.com>

Abstract

This report examines the current state-of-the-art in large-scale system configuration. It introduces the scope and principles of the configuration task, and describes a selected range of technologies that are representative of approaches in use today. We conclude by summarising the dominant approaches to system configuration offered by these technologies, and by indicating requirements that are currently unmet — which will lead on to future research.

The report is a deliverable of the GridWeaver project, a collaboration between the University of Edinburgh and HP Laboratories in Bristol. GridWeaver is part of the UK e-Science Core Programme.

Revision 0.2 – November 12, 2002

Next term

- Following the workplan:
 - Future requirements
 - Design of proof-of-concept implementations
 - Implementation of pilot software
- Working hard to finish the first two by end Jan-03
 - Aim to safeguard 6 months of work on pilot software
- Overlap of tasks helps:
 - Better skill and personnel utilisation
 - Important benefits from looking ahead

Future requirements

- D2.1: Public report expressing our expectations for the Grid fabric requirements
 - We will also highlight the problems that we will investigate in the rest of the project
- The bulk will be extracted from *case studies* of large installations around the world
 - Interviews of fabric managers in 4-5 organisations
 - Academic
 - Research
 - Industrial
- We also plan to include *hypothetical case studies* to found our beliefs for pitfalls not experienced yet
 - ...but will have to be negotiated sooner or later

Proof-of-concept implementations

- D3.1: Design of the software deliverable due next
 - We want to combine:
 - The descriptive power of SmartFrog, and its focus on cross-node, distributed services
 - The proven abilities of LCFG
 - We aim to extract the resulting *configuration model*

- D4.1: Software to experiment on
 - Investigate the configuration model
 - Evaluate the abilities of the SmartFrog language
 - No production code, plenty of sticky tape
 - Real life scenarios (firewall configuration?)

- D4.2: Final report with our observations

After this project (1)

- Collaboration is proceeding very well
 - Common vision for large-scale configuration
 - Shared view on the problems
 - Complementary technologies and perspectives
 - LCFG, SmartFrog
 - Academia, industry
 - Enthusiasm, insight, experience, engagement and momentum from the team members
- Initial 1 year project will only scratch the surface
 - Project intended as a “trial” engagement
 - Will define additional research problems that are not within the scope of a small, 1 year project
 - Will develop proof-of-concept implementations, but not releasable systems

After this project (2)

- Numerous opportunities for further research and system development, e.g.:
 - Robust, “universal”, declarative configuration specification language
 - templates for common fabric configurations
 - composable predicates for automatic correctness checking
 - ...
 - Range of models and control components for commonly used fabric resources
 - how to model and automatically configure/control a diversity of resources, at various levels of abstraction

Conclusions

- The project addresses fundamental issues
 - No fabric-no applications
- Partners working well together
 - Common mentality, complementary skills
- Good progress with work so far
 - First deliverable can prove useful to a variety of audience
 - Realistic, achievable, well-monitored schedule
- Lots of areas for future work
 - This project merely sets the scene